IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Applicant:	Lolayekar, Santosh C. et al.		
Serial No.	10/051,339	Group Art Unit:	2154
Filed:	January 18, 2002	Examiner:	Lin, Wen Tai
Title:	Enforcing Quality of Service in a Storage Network		

APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

Dear Sir:

Enclosed is Appellants' Appeal Brief pursuant to 37 C.F.R. § 41.37 in connection with the Notice of Appeal filed April 28, 2006 from the final rejection of Claims 1, 3-25 and 27-44 in the Office Action of January 31, 2006 ("Office Action"). Also enclosed is a Petition for Extension of Time for a one (1) month extension. The fees for the Appeal Brief and the Extension of Time for a Large Entity are enclosed.

TABLE OF CONTENTS

TABLE OF AUTHORITIES

Cases

Statutes

Other Authorities

I. REAL PARTY IN INTEREST

The real party in interest is EMC Corporation, a corporation of the Commonwealth of Massachusetts. This application was recently assigned to EMC Corporation, a Large Entity. Accordingly, the claim for Small Entity Status is hereby withdrawn.

II. RELATED APPEALS AND INTERFERENCES

There are no pending appeals, interferences or judicial proceedings known to Appellants, to Appellants' legal representative, or to Assignee which may be related to, directly effect, be effected by, or have a bearing on the Board's decision in this appeal.

III. STATUS OF CLAIMS

Claims 1, 3-25 and 27-44 are pending; Claims 1, 3-25 and 27-44 stand finally rejected by the Examiner; and Claims 1, 3-25 and 27-44 are being appealed.

IV. STATUS OF AMENDMENTS

A Response to the Final Rejection was filed March 31, 2006. An Advisory Action dated April 17, 2006 was issued maintaining the final rejection of Claims 1, 3-25 and 27-44, but indicating the Response of March 31, 2006 would be entered for purposes of appeal.

V. SUMMARY OF CLAIMED SUBJECT MATTER

A. Concise Explanation of Subject Matter Claimed

The invention relates to methods for use in storage networks, such as storage area networks ("SANs"), to intelligent storage switches for use in storage networks, and to machine readable media embodying instructions for execution by a storage switch in a storage network for affording a Quality of Service (QoS) connection between initiators, such as servers, and targets, such as storage devices. The invention is described in the specification generally in the Summary, paragraphs [0013] – [0016], and in paragraphs [0018 – [0024] and [0029] – [0035], and shown in Figures 2-5.

A storage switch in accordance with the invention enables centralized management of globally distributed storage devices which may be used in shared storage pools, while avoiding the necessity of the large number of globally distributed management stations and skilled management personnel which characterize conventional storage networks. Unlike conventional SAN switches which have little built-in intelligence and merely forward data to a selected appliance in the storage network that provides storage management (specification, paragraph [0009]), storage switches in accordance with the invention distribute intelligence to every switch port and provide, in addition to switching functions, virtualization, storage services and other services such as Quality of Service for storage access that are typically performed by separate appliances and other devices in conventional storage networks (paragraph [0024]).

Figure 5 illustrates a functional block diagram of a storage switch 204 in accordance with the invention. As shown, the switch may comprise a plurality of linecards 502, 504, and 506, a plurality of fabric cards 508, and system control cards 510. (See paragraph [0029]). The system control cards may monitor the individual linecards and the fabric cards, and maintain a database 512 that tracks configuration information of virtual targets and physical devices such as servers and storage devices that are attached to the switch (paragraph [0031]). The fabric cards interface to the switching matrix for the storage switch and to the linecards. The linecards connect to the servers and to the storage devices via ports 602 (Figure 6, paragraph [0035]).

Figure 6 is a functional block diagram of an embodiment of a generic linecard of a storage switch 204 in accordance with the invention. As shown, the linecard may comprise a plurality of ports 602, eight ports being shown in the figure. The ports are full duplex and connect to either a server or other client, or to a storage device or subsystem as shown, for example, in Figures 2-5. (See paragraph [0037]). Each port has an associated memory 603 and an associated Storage Processor Unit (SPU) 601. The SPU may include several elements, a Packet Aggregation and Classification Engine (PACE) 604, a Packet Processing Unit (PPU) 606, an SRAM 605, and a CAM 607 (paragraph [0039]). Each PACE element aggregates two ports into a single data channel having twice the bandwidth (paragraph [0040]), and adds local headers to data packets and sends the data packets to a Packet Processing Unit (PPU) 606. Each PPU is divided into an ingress PPU and an egress PPU, and

performs virtualization and protocol translation on-the-fly, meaning that cells are not buffered. (See paragraph [0042]).

Each linecard 600 may also include two traffic managers (TMs) 608, one for ingress traffic and one for egress traffic. The ingress TM receives data cells from the SPUs and sends the data cells to the fabric cards 610, and the egress TM receives data cells from the fabric cards and delivers them to the SPUs. (See paragraphs [0048] – [0049]). The fabric cards connect to the fabric switching matrix.

As further shown in Figure 6, each linecard also includes a processor (CPU) 614 which is responsible for initializing every chip on the card at power up and for processing control traffic during operation.

A storage switch in accordance with the invention may perform various switch-based storage operations, including pooling and provisioning, Quality of Service ("QoS") for storage access, and load balancing. (See paragraph [0066]). A storage switch in accordance with the invention enforces QoS by guaranteeing a minimum percentage of bandwidth for an initiator connection to a storage device. The switch may determine the data bandwidth by calculating the number of requests per second and multiplying the requests by the average transfer size of the request (See paragraph [0109]). The traffic managers 608 monitor the transfer bandwidth of the different connections, and schedule delivery based upon QoS parameters. The TMs guarantee that each shared connection gets its appropriate minimum bandwidth, and may accumulate packets in a TM buffer 612 in order to limit the bandwidth to within a specified range. As packets accumulate in the buffer, the traffic manager

may send a message to the SPU to inform an initiator to slow its connection. The switch will also match the bandwidth between an initiator and a storage device by allocating an appropriate proportion of the bandwidth.

B. Correspondence Between the Claims and the Specification

The following indicates (in bold) the correspondence between the specification paragraphs (in square brackets) and reference characters of the drawings for subject matter defined by the independent claims on appeal.

1. Independent Claims

1. A method for use in a storage network **(200 (Fig. 2), 302, 304 (Fig.'s 3-4, 8-11, 13))**, the storage network including at least one initiator **(server 202 (Fig.'s 2-4), [0018])** at least one storage device **(206, 207 (Fig.'s 2-4), [0019])**, and a storage switch **(204 (Fig's 2-5, 8-11, 13), [0015] - [0018])** in communication with the at least one initiator and the at least one storage device **([0018] – [0019])**, the method comprising:

providing, by the storage switch, quality of service to the at least one initiator for accessing the at least one storage device in the storage network **([0024], [0066])**.

9. A method for use in a storage network **(200 (Fig. 2), 302, 304 (Fig.'s 3-4, 8-11, 13))**, the storage network including at least one initiator **(server 202 (Fig.'s 2-4), [0018])**, at least one storage device **(206, 207 (Fig.'s 2-4), [0019])**, and at least one storage switch **(204 (Fig's 2-5, 8-11, 13), [0015] - [0018])**, wherein the at least one initiator and the at least one storage device are both in communication with the storage switch **(204 (Fig.'s 2-5, 8-11, 13), [0015] - [0018])**, the method comprising:

guaranteeing, by the storage switch, a minimum bandwidth to the at least one initiator to access the at least one storage device in the storage network **([0107])**; and

measuring, by the storage switch, an actual bandwidth utilized by the at least one initiator, where the actual bandwidth is measured by a number of requests per

second times an average size of requests from the at least one initiator **([0069], [0107])**.

16. A method for use in a storage network **(200 (Fig. 2), 302, 304 (Fig.'s 3-4, 8-11, 13))**, the storage network including a plurality of initiators **(servers 202 (Fig.'s 2-4), [0018])**, a plurality of targets **(storage devices 206, 207 (Fig.'s 2-4), [0019])**, and a storage switch **(204 (Fig.'s 2-5, 8-11, 13), [0015] - [0018])**, the method comprising:

guaranteeing, by the storage switch, a respective minimum bandwidth for each of a plurality of connections **([0107], [0115])**, wherein each respective connection is a connection from a respective initiator to a respective target via the storage switch in the storage network **(Fig.'s 2-5)**;

monitoring, by the storage switch, an actual bandwidth utilized by each initiator **([0111])**, where the actual bandwidth is measured by a number of requests per second from the initiator times an average size of the requests from the initiator **([0069], [0109])**; and

determining if the actual bandwidth used by one initiator is excessive, and, if excessive, adjusting, by the storage switch, a number of allowed concurrent requests or at least one initiator **([0110], [0111])**.

22. A method for use in a storage network **(200 (Fig. 2), 302, 304 (Fig.'s 3-4, 8-11, 13))**, the storage network including at least one initiator **(server 202 (Fig.'s 2-4), [0018])**, at least one storage device **(206, 207 (Fig.'s 2-4), [0019])**, and a storage switch **(204 (Fig.'s 2-5, 8-11, 13), [0015] - [0018])**, wherein the at least one initiator and the at least one storage device are both in communication with the storage switch **(Fig.'s 2-5)**, the method comprising:

providing a connection from the at least one initiator to the at least one storage device via the storage switch in the storage network **(204 (Fig.'s 2-5, 8-11, 13), [0015] - [0018])**; and

adjusting, by the storage switch, the number of requests allowed the at least one initiator to keep the bandwidth utilized by the at least one initiator within a specified range **([0109], [0111])**.

25. A switch **(204 (Fig.'s 2-5, 8-11, 13), [0015]-[0018])** for use in a storage network **(200 (Fig. 2), 302, 304 (Fig.'s 3-4, 8-11, 13))**, the switch comprising:

a port **(602, Fig. 6, [0037])** to be coupled to an external device, wherein the external device includes at least one of an initiator and a storage device **([0037])**; and

a bandwidth controller **([0111])**, the bandwidth controller including a processor **(SPU 601, Fig. 6, [0111])**, a traffic manager **(TM 608, Fig. 6, [0111])**, and a buffer **(612, Fig. 6, [0111])**.

30. A switch **(204 (Fig.'s 2-5, 8-11, 13), [0015] - [0018])**, including:

a storage processor **(SPU 601, Fig. 6, [0111])**, including a request controller;

a traffic manager **(TM 608, Fig. 6, [0111])** in communication with the storage processor;

a buffer **(612, Fig. 6, [0111])** in communication with the traffic manager;

wherein if a specified threshold in the buffer is reached, the traffic manager is designed to activate the request controller **([0111])**.

37. A storage network **(200 (Fig. 2), 302, 304 (Fig.'s 3-4, 8-11, 13))**, including:

an initiator **(server 202 (Fig.'s 2-4), [0018])**;

a storage device **(206, 207 (Fig.'s 2-4), [0019])**;

a switch **(204 (Fig.'s 2-5, 8-11, 13), [0015] - [0018])** in communication with the initiator and the storage device **(Fig.'s 2-5)**;

wherein the switch includes a traffic manager **(TM 608, Fig. 6, [0111])** in communication with a buffer **(612, Fig. 6, [0111])**;

wherein when the buffer includes a number of packets from the initiator that exceeds a specified threshold, then the switch is designed to notify the initiator to reduce a number of concurrent requests **([0109], [0110])**.

38. A machine readable media which has instructions stored thereon, which when executed by a storage switch **(204 (Fig.'s 2-5, 8-11, 13), [0015] - [0018])** in a storage network (**200 (Fig. 2), 302, 304 (Fig.'s 3-4, 8-11, 13)** including an initiator **(server 202 (Fig.'s 2-4), [0018])** and a storage device **(206, 207 (Fig.'s 2-4), [0019])** in communication with the storage switch causes the storage switch to perform the following steps:

guaranteeing, by the storage switch, a minimum bandwidth to the initiator to access the storage device in the storage network **([0107], [0115])**; and

measuring, by the storage switch, an actual bandwidth utilized by the initiator, where the actual bandwidth is measured by a number of requests per second times an average size of requests from the initiator **([0069], [0107])**.

2. Means-Plus-Function Claims

The following indicates in bold the correspondence between the specification paragraphs and reference characters of the drawings for each means-plus-function element of the independent and dependent claims on appeal.

33. A storage switch **(204 (Fig.'s 2-5, 8-11, 13), [0015] - [0018])** for use in a storage network (**200 (Fig. 2), 302, 304 (Fig.'s 3-4, 8-11, 13)**) comprising:

a first port **(602 (Fig. 6), [0037])** to be coupled to at least one initiator **(202)**;

a second port **(602 (Fig. 6), [0037])** to be coupled to at least one storage device **(206, 207)**;

means **(linecard 600, TM 605, SPU 601 (Fig. 6))** for providing quality of service for a connection from the at least one initiator to the at least one storage device in the storage network.

34. The switch of claim 33, wherein means for providing quality of service includes:

means **(TM 608, [0111])** for guaranteeing a minimum bandwidth to the at least one initiator to access a storage device;

means **(TM 608, [0112])** for measuring an actual bandwidth utilized by the at least one initiator, where the actual bandwidth is measured by the number of requests per second times the average size of the requests from the at least initiator; and

means **(SPU 601, [0112])** for adjusting the number of concurrent requests allowed to be sent by the at least one initiator to keep the bandwidth utilized by the at least one initiator within a specified range having as a lower limit the minimum bandwidth.

35. The switch of claim 33, wherein means for providing quality of service includes:

a processor **(SPU 601, [0039], Fig. 6, [0111] - [0112]);**
a traffic manager **(TM 608, [0048], [0111], Fig. 6);** and
a buffer **(612, [0111], Fig. 6)**.

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Whether Claims 1, 25, 27, 30, 33, and 35-37 are unpatentable under 35 U.S.C. §102(e) as anticipated by Published Application U.S. 2002/0194324 to Guha ("Guha") for the reasons stated in paragraphs 3-7 and 23 of the Office Action.

2. Whether Claims 3-24, 29, 31-32, 34 and 38-44 are unpatentable under 35 U.S.C. §103(a) as obvious over Guha (U.S. 2002/0194324) in view of U.S. Patent No. 5,719,854 to Choudhury et al. ("Choudhury") for the reasons stated in paragraphs 8-20 of the Office Action.

VII. ARGUMENT

For the reasons set forth below, it is respectfully submitted that the grounds of rejection of the claims are legally in error, substantively and procedurally, factually unsustainable, and should be reversed.

A. The Rejections of the Claims Are Improper Because the Office Legally Misconstrued the Claims

The Office's rejection of Claims 1, 25, 27, 30, 33 and 35-37 under 35 U.S.C. §102(e) as anticipated by Published Application U.S. 2002/0194324 to Guha, and its rejections of Claims 3-24, 29, 31-31, 34 and 38-44 under 35 U.S.C. §103 as unpatentable over Guha in view of U.S. 5,719,854 to Choudhury et al. are improper because the Office's claim construction is legally incorrect.

It is well settled that the first step in evaluating claims for patentability is to construe the claims to determine their claim scope. (See *In re Hiniker Co.*, 150 F.3d 1362, 1369, 47 U.S.P.Q.2d 1523, 1529 (Fed. Cir. 1998)). Here, the Office has improperly construed the claim term "storage switch", a physical structural element (a device), as corresponding to an "entity" comprising a plurality of separate, discrete physical elements disclosed in Guha. The Office asserts as the basis for this construction that the "entity" comprising the separate devices of Guha collectively incorporates the functionalities of Applicants' claimed storage switch.

It is respectfully submitted that the Office's claim construction is legally wrong and contrary to well-established legal authority, and that its rejection of the claims based upon this erroneous claim construction is improper.

1. Claim Terms Must be Interpreted as They Would be by One of Ordinary Skill in the Art and Consistent With the Specification

There is a large body of well-established legal precedent which governs the construction of claims. Some of this precedent is summarized in M.P.E.P. §2111, which provides guidance to Examiners in construing claims.

During prosecution, Examiners give claims their broadest reasonable interpretation. However, this interpretation is not unbounded. Claim terms must be construed to have the meanings that the one that those skilled in the art would reach. See *In re Cortright*, 165 F.3d 1353, 1358, 49 U.S.P.Q.2d 1464, 1468 (Fed. Cir. 1999); *In re Morris*, 127 F.3d 1048, 1054, 1056, 44 U.S.P.Q.2d 1023, 1029 (Fed. Cir. 1997) ("the PTO applies to the verbiage of the proposed claims the broadest reasonable meaning of the words in their ordinary usage as they would be understood by one of ordinary skill in the art"); *In re Bond*, 910 F.2d 831, 833, 15 U.S.P.Q.2d 1566, 1567 (Fed. Cir. 1990) ("it is axiomatic that, in proceedings before the PTO, claims in an application are to be given their broadest reasonable interpretation consistent with the specification and that the claim language should be read in light of the specification as it would be interpreted by one of ordinary skill in the art"); M.P.E.P. §2111.01 ("the words of a claim must be read as they would be interpreted by those of ordinary skill in the art"). Prior art references may be "indicative of what all those skilled in the art generally believe a certain term means ... [and] can often help to demonstrate how a disputed term is used by those skilled in the art", *Vitronics Corp. v. Conceptronic, Inc.*, 90 F.3d 1576, 1584, 39 U.S.P.Q.2d 1573, 1578-79 (Fed. Cir. 1996). Accordingly, the PTO's interpretation of claim terms should not conflict with the

meanings given to identical terms in the specification, by those skilled in the art, and in other patents from analogous art.

Here, the Office has failed to correctly apply these legal standards, and has misconstrued the claim term "storage switch" by interpreting that term to read on a plurality of separate and discrete physical devices, contrary to Applicants' own specification, contrary to the understanding of term storage switch to those of ordinary skill of the art, and even contrary to the use of the term switch in the cited prior art reference to Guha upon which the Office bases its rejections.

2. The Specification Uses the Term "Storage Switch" According to the meaning of That Term as Understood by One Skilled in the Art

The invention relates to methods and storage switches employed in storage networks, and in particular in storage area networks (SANs). The specification describes and illustrates a conventional storage area network on pages 3-5 and Figure 1, as comprising a plurality of servers 102 connected to a plurality of storage devices 116 through Fibre Channel switches 112, and appliances 114. The Fibre Channel switches and appliances are controlled by storage managers 120, and are connected to gateways 118 and external networks via a router 108 and bridges 121 to provide storage services.

As described in Applicants' specification, a key aspect of the invention is to replace a conventional storage switch, such as the Fibre Channel switch 112 shown in Figure 1, in a storage network with the intelligent storage switch of the invention that performs some functions previously performed by the separate Fibre Channel switches, appliances and gateways in the conventional storage network, and to

incorporate into the intelligent storage switch additional functions, such as Quality of Service performed in conventional systems by other separate devices. (See paragraphs [0024] - [0026]). The Office improperly justifies its claim construction of "switch" as corresponding to a plurality of separate discrete devices on the basis that the specification describes the additional functions performed by the switch of the invention as being performed by other devices in conventional systems (see, Office Action, page 8, paragraph 23). The Office's construction is not only legally incorrect, it is illogical.

The specification describes the storage switch of the invention as a unitary structure, i.e., a physical thing, a device, and emphasizes the enhanced functionality of the improved storage switch of the invention over conventional Fibre Channel storage switches by describing functions of the switch of the invention that were not previously performed by the prior art switches. The specification clearly describes and uses the term "storage switch" in reference to the invention to refer to a thing or device, which is the same way it describes the conventional Fibre Channel switch, not to a collection of devices as construed by the Office. This use of the term "storage switch" in the specification as meaning a device is consistent with the meaning of the term as it is understood by those skilled in the art.

3. Those Skilled in the Art Understand the Term "Storage Switch" to Mean a Device

The plain ordinary meaning of the term "switch" is a thing or device. As understood by those skilled in the art, the term "storage switch" means a thing or device that interconnects servers and storage devices in a storage network and

performs switching functions. Evidence of this meaning and understanding is demonstrated by Exhibits A – D in the Evidence Appendix. These Exhibits, which were previously submitted in Applicants' Response of March 31, 2006, are clear evidence of the meaning of the term "switch" to one skilled in the art.

Exhibits A and B are white papers that describe storage networks and switches, and clearly describe a storage switch as being a device. Exhibit A describes SANs as using special switches which "look a lot like a normal Ethernet networking switch" as a mechanism for connecting devices (See page 2). Exhibit B compares a Fibre Channel switch to an Ethernet switch and states that it performs the same basic functions as a switch on an Ethernet network in that it acts as a connectivity point for devices (See page 2). Exhibit C defines "storage switch" as a "device that routes data between servers and disk arrays in a storage area network (SAN). It typically refers to a Fibre Channel switch". Finally, Exhibit D defines a "Fibre Channel switch as a computer storage device that allows the creation of a Fibre Channel fabric".

These Exhibits clearly demonstrate that the term "storage switch" is understood by those skilled in the art to comprise a device, i.e., a physical element. The term is clearly not understood by those skilled in the art to comprise a collection of separate, discrete devices which happen to perform some similar functions, as asserted by the Office in its rejection.

4. The Prior Art Relied Upon by the Office Uses the Term "Switch" to Mean a Device

Guha, the primary reference relied upon by the Office for its rejections, also uses the term "switch" to mean a device. Guha illustrates in Figures 4 and 6 and describes conventional storage networks as comprising a plurality of separate devices including servers 39, 40 and 41 connected to storage devices 44 by a SAN switch 42. Guha describes the QoS enforcer 34 as comprising a "network routing device, preferably a load balancing network device" (Guha, paragraph [0048]). Guha further describes that the "QoS enforcer 34 communicates through a layer switch 38 with a plurality of servers, such as a large content server 39, a web server 41 and a database server 40" and that the "servers communicate with storage devices 44 through the network storage switch 42 (i.e., SAN switch)". (See Guha, paragraph [0063]). Thus, Guha describes the SAN switch 42 as being a device which connects servers to storage devices in a storage network. SAN switch 42 of Guha is analogous to the storage switch of the invention which connects initiators (*i.e.*, servers) to storage devices. Accordingly, Guha is further evidence of the understanding of the term "storage switch" to one skilled in the art as being a device and not an "entity" comprising a collection of separate discrete entities, as asserted by the Office.

B. The Rejection of Claims 1, 25, 27, 30, 33 and 35-37 as Anticipated Under 35 U.S.C. §102 is Improper

Claims 1, 25, 27, 30, 33 and 35-37 cannot be anticipated by Guha, and the rejection of these Claims under 35 U.S.C. §102(e) is improper and should be reversed.

In order for the prior art to anticipate under 35 U.S.C. §102, every element of a claimed invention must be identically shown in a single prior art reference, and all elements in the reference must be arranged as in the claim under review. (See *In re Bond, supra*). Guha does not identically disclose all of the elements of any of Claims 1, 25, 27, 30, 33 or 35-37.

1. Claim 1

Independent Claim 1 recites:

> "A method for use in a storage network, the storage network including at least one initiator, at least one storage device, and a storage switch in communication with the at least one initiator and the at least one storage device, the method comprising:
>
> providing, by the storage switch, quality of service to the at least one initiator for accessing the at least one storage device in the storage network."

As demonstrated above, the Office has improperly construed the claims by reading the term "storage switch" as comprising the content controller 36, the QoS enforcer 34, the Layer 4 switch 38, and the SAN switch 42 of Guha, which the Office asserts together constitute "an entity" that corresponds to Applicants' "storage switch" (see Office Action page 2, paragraph 5 and page 8, paragraph 23). This interpretation is contrary to Applicants' specification, contrary to the understanding of the term storage switch to those skilled in the art, and contrary to the meaning of the term switch as used in the Guha reference, for the reasons discussed above. The rejections based upon this construction constitute legal error, and should be reversed.

Guha does not disclose a storage switch which provides Quality of Service to an initiator for accessing a storage device, as recited in Claim 1. Rather, Guha

discloses a Quality of Service enforcer 34, separate from the SAN switch 42, which he explicitly describes as being a network routing device, preferably a load balancing device (emphasis added)(see paragraph, [0048]). Moreover, Guha discloses that the QoS enforcer prioritizes content requests from an external network to the application servers (emphasis added)(see paragraph [0068]). Guha does not disclose a QoS enforcer that provides quality of service to an initiator for accessing a storage device, as claimed. Accordingly, Guha cannot anticipate Claim 1.

2. Claims 25 and 27

Independent Claim 25 is directed to a switch for use in a storage network that comprises:

> "a port to be coupled to an external device comprising at least one of an initiator and a storage device, and
> a bandwidth controller including a processor, a traffic manager and a buffer"

Dependent Claim 27 calls for the processor to be a storage processor.

Guha does not disclose a switch that embodies any of the elements set forth in Claims 25 and 27. For the reasons explained above, the separate devices of Guha cannot constitute an "entity" corresponding to a storage switch, as claimed. Rather, Guha discloses a SAN switch 42, which is a generic storage switch of the type to which the invention relates, and a plurality of devices that are separate and discrete from the SAN switch, which the Office asserts (improperly) together comprise the claimed storage switch.

In its rejection of Claims 25 and 27, the Office asserts (see Office Action page 3, paragraph 6) that the claimed port coupled to an external device (that includes one of an initiator and a storage device) comprises element 64 in Figure 6 of Guha. However, Guha discloses element 64 to be a router which routes a request from an external Internet through load balancer 35, the QoS enforcer 34, the layer switch 38, servers 39-41 and SAN switch 42 to the storage system 58 (see paragraph [0067]). Thus, contrary to the Office, router 64 is not a port of the switch, as claimed.

The Office further asserts (Office Action, paragraph 6) with regard to the claimed bandwidth controller that there inherently must be a processor in the SAN switch, that the traffic manager corresponds to element 34 (QoS enforcer) or 36 (content controller), and "that a buffer must have existed in order to perform the tasks described at paragraphs 68 - 69". The Office's position on these elements comprises pure speculation for which there is no support in Guha.

Guha does not disclose a bandwidth controller at all, much less a bandwidth controller which includes a processor, a traffic manager and a buffer, as claimed. The Office has not presented any reasonable or logical basis for its assertions that Guha inherently includes the claimed elements. Rather, in contrast, Guha, at paragraphs [0068] – [0069] and in Figure 6, describes that the QoS enforcer 34 and the content controller 36 prioritize and delay or drop content requests input from the Internet through the router to the application servers in order to satisfy service level agreements (SLA). This has nothing to do with either providing quality of service by a storage switch to an initiator for accessing a storage controller (Claim 1), or for controlling bandwidth, as claimed in Claims 25. Accordingly, it is respectfully

submitted that the rejection of Claim 25 as anticipated by Guha is improper and should be reversed.

Claim 27, which sets forth that the processor is a storage processor depends from Claim 25, and is deemed to be allowable over Guha for at least the same reasons that Claim 25 is allowable. Moreover, even assuming, *arguendo,* that the SAN switch of Guha includes a processor as asserted, there is no disclosure in Guha of the processor being a storage processor, as claimed.

3. Claims 30, 33 and 35-37

Claims 30, 33 and 37 are separate independent claims that recite different elements. The Office's lumping together and rejecting these claims as being anticipated by Guha (see Office Action page 3, paragraph 7) with the broad, sweeping generalization that the features of these Claims can be found in Claims 1, 25 and 27 and that they are rejected for the same reasons is wholly inadequate to satisfy its requirements for establishing a *prima facie* rejection for anticipation. To satisfy its burden of establishing a *prima facie* rejection, the Office must demonstrate identity between the elements of each claim and the Guha reference (see *In re Bond, supra*). This, the Office has not done, and cannot do, since Guha does not identically disclose all of the elements of any of the independent claims.

Independent Claim 30 is directed to a switch that includes a storage processor including a request controller, a traffic manager in communication with the storage processor, and a buffer in communication with the traffic manager, and whereupon a specified threshold in the buffer being reached, the traffic manager activates the

request controller. There is no disclosure (or suggestion) in Guha of a switch including the elements set forth in Claim 30 either in a single device or in multiple devices, and Guha cannot anticipate Claim 30 for this reason.

Claim 33 calls for a storage switch comprising:

> "a first port to be coupled to at least one initiator;
> a second port to be coupled to at least one storage device, and
> means for providing Quality of Service for a connection from the at least one initiator to the at least one storage device in the storage network."

Independent Claim 33 clearly contemplates that the storage switch is a device having first and second ports, not a collection of independent devices as asserted by the Office. Guha does not disclose a switch comprising a device having first and second ports as set forth in Claim 33.

a. Claim 33 Must be Construed Pursuant to 35 U.S.C. §112, ¶6

Moreover, the Office has improperly construed Claim 33 for another reason. The last element of Claim 33 is written in means-plus-function language, and must be interpreted in accordance with 35 U.S.C. §112, ¶6, which requires the language to be construed as covering the corresponding structure, material or acts described in the specification for performing the recited function and equivalents thereof. The Office must interpret means-plus-function language in claims in accordance with the statute during prosecution, and therefore, is required to consult the specification in order to determine the permissible scope of the claim. (See *In re Donaldson*, 16 F.3d 1189; 29 U.S.P.Q.2d 1845 (Fed. Cir. 1994)). The Office must construe the means-plus-function language of Claim 33 as covering the corresponding structure described in

the specification for performing the recited function, and it must find that structure or equivalent structure in Guha in order to satisfy its burden of establishing anticipation. The Office has failed to do this, and the rejection of Claim 33, as well as Claim 34 which depends thereon, cannot be sustained.

Finally, independent Claim 37 is directed to a storage network including an initiator, a storage device, and a switch in communication with the initiator and storage device, where the switch includes a traffic manager in connection with a buffer, and when the buffer includes a number of packets from the initiator that exceed a specified threshold, the switch notifies the initiator to reduce the number of concurrent requests.

There is no disclosure (or suggestion) in Guha of a storage network that includes a switch comprising the elements set forth in the claim and which, upon the number of packets in a buffer exceeding a specified threshold, notifies an initiator to reduce the number of concurrent requests, as claimed. Accordingly, Guha cannot anticipate Claim 37.

In view of the foregoing, it is respectfully submitted that the rejection of Claims 1, 25, 27, 30, 33 and 35-37 under 35 U.S.C. §102(e) as anticipated by Guha is improper and should be reversed. Furthermore, since Guha neither teaches nor suggests a storage switch and a storage network that includes the elements recited in these claims, it is also submitted that Guha cannot render these claims obvious.

C. The Rejections Under 35 U.S.C. §103 Are Improper and Should Be Reversed

The rejections of the claims under 35 U.S.C. §103 as obvious over Guha in view of Choudhury are improper and should be reversed. Not only can Guha and Choudhury not be combined in the manner stated in the Office Action to support the rejection, even if the references were combined, the resulting combination would not produce the claimed invention.

In order to combine references to support a rejection under Section 103, there must be some objective teaching in the references themselves or in the knowledge generally available to one of ordinary skill in the art that would lead that individual to combine the relevant teachings of the references. See *In re Kahn,* 441 F.3d 977, 78 U.S.P.Q.2d 1329 (Fed. Cir. 2006); *In re Fine*, 837 F.2d 1071, 5 U.S.P.Q.2d 1596 (Fed. Cir. 1988). Even where all aspects of a claimed invention were individually known in the art, there must be some motivation, suggestion, or teaching of the desirability of making the specific combination claimed. In the absence of a showing of such a motivation to combine the teachings of references, no *prima facie* case of obviousness can be made. See *In re Kotzab*, 217 F.3d 1365, 55 U.S.P.Q.2d 1313 (Fed. Cir. 2000). Here, it is respectfully submitted that there is no motivation, suggestion or teaching of the desirability of combining Guha and Choudhury, as asserted by the Office.

As discussed above, Guha discloses a storage area network that comprises a plurality of servers that are connected to a plurality of storage devices by a SAN switch, and where the system further includes a content controller and a QoS

enforcer that directs external requests to servers via a load balancer, a Layer 4 switch, to enforce SLAs with the servers. Although Guha and the present invention may both be directed to storage networks, they address different problems and use far different approaches.

Guha discloses a conventional storage area network to which has been added separate appliances for performing the functions of Quality of Service enforcer and content controller. In contrast, the invention provides Quality of Service in a storage area network for requests from an initiator (server) to a storage device through a novel intelligent storage switch that replaces a conventional storage switch, such as a SAN switch as disclosed by Guha. There is no disclosure or suggestion anywhere in Guha of providing Quality of Service between initiators and storage devices in a storage network by incorporating functionality for that purpose in an intelligent storage switch, as claimed. Accordingly, Guha cannot render any of the claims obvious.

Choudhury, et al. does not cure the deficiencies of Guha, and cannot in combination with Guha render the claims obvious. Choudhury discloses methods for providing multiple grades of service with protection against overloads for multiple customers sharing limited resources in connection with systems such as circuit-switched telecommunication networks and broadband integrated-services digital networks. Choudhury does not disclose or discuss storage networks. The disclosed method controls access to the networks by multiple customers using streams that regulate the requests admitted from different customers and by providing grades of services to customers that share a resource. Each customer is provided an upper limit (UL) and a guaranteed minimum (GM) on the number of requests from that

customer that can put into service at any time. (column 6, lines 21-28). The reference discloses solving a resource-sharing model with the grades of services in effect for customers using a "blocking probability computer" process (column 7, lines 22-27) to determine in real time whether a new perspective customer can be admitted to a resource with a desired grade of service. If all blocking requirements can be met, then the resource provider admits the new customer with the desired grade of service. However, if all blocking requirements cannot be met, then the blocking probability computer is used to determine whether a lower grade of service is feasible, and if so, the lower grade is offered to the customer (column 8, lines 36-51).

D. Claims 3-24, 29, 31-32, 34 and 38-44 Would Not Have Been Obvious Over Guha and Choudhury

1. Claims 3-5

Dependent Claim 3, which depends from Claim 1, recites that the step of providing Quality of Service by the storage switch includes controlling the number of packets from the at least one initiator to the at least one storage device during a period of time. Dependent Claim 4 calls for controlling the number of requests from the at least one initiator to the at least one storage device; and dependent Claim 5 calls for adjusting a number of concurrent requests allowed to be sent by the at least one initiator.

In rejecting Claims 3-5 (Office Action, page 4, paragraph 10), the Office acknowledges that Guha does not specifically teach controlling the number of packets or concurrent requests from an initiator to a storage device. It asserts, however, that Choudhury teaches controlling the number of concurrent requests sent

from an initiator as an indication of the network traffic caused by the initiator (referring to the Abstract of Choudhury), and further asserts that it would have been obvious to use the number of packets or concurrent requests sent out from an initiator as a measure of traffic load and determine whether the request would be granted. (See Office Action, page 4, paragraph 10).

Even assuming, *arguendo,* the correctness of the Office's assertion, this is not what is claimed in Claims 3-5. In particular, Claim 3 requires controlling the number of packets from an initiator to a storage device during a period of time; Claim 4 requires controlling the number of requests from the initiator to the storage device; and Claim 5 requires adjusting a number of concurrent requests that are allowed to be sent by the initiator.

Contrary to the Office's assertion, Choudhury does not control or adjust either packets or requests for service. Rather, Choudhury only discloses whether to connect or deny service to a requestor for a given grade of service (see column 9, line 61-column 10, line 15). Nothing in Choudhury teaches or suggests controlling either the number of packets or the number of requests from an initiator to a storage device, nor adjusting a number of concurrent requests that are allowed to be sent by the initiator, as claimed. Accordingly, the rejection of Claims 3-5 on Guha and Choudhury is improper, and should be reversed.

2. Claims 6-8

Dependent Claim 6, which depends from Claim 1, sets forth that the step of providing Quality of Service includes adjusting the number of requests allowed the at

least one initiator to keep the bandwidth utilized by the one initiator within a specified range.

Dependent Claim 7 sets forth at the step of providing Quality of Service includes guaranteeing a minimum bandwidth to the at least one initiator to access the at least one storage device; measuring an actual bandwidth utilized, and adjusting a number of concurrent requests allowed to be sent by the initiator.

Dependent Claim 8 calls for guaranteeing up to a maximum bandwidth to the at least one initiator to access the at least one storage device; where adjusting the number of concurrent requests includes reducing the number of concurrent requests allowed by the one initiator when the actual bandwidth exceeds the maximum bandwidth.

In its rejection of Claims 6-8, the Office asserts (Office Action, page 5, paragraph 11) that while Guha does not teach adjusting the number of requests allowed to keep the bandwidth within a specified range, Choudhury teaches that resource capacity may be measured as available bandwidth and that bandwidth may be measured by the number of requests, and that it would have been obvious to adjust an initiator's use of network bandwidth to within a specified range. This is incorrect. At column 14, lines 9-16, referred to by the Office, Choudhury teaches only determining the cost of service to a customer in terms of capacity between a guaranteed minimum (GM) and a maximum upper limit (UL). Accordingly, Choudhury will not support the rejections.

As discussed above, neither reference discloses or suggests controlling or adjusting the number of requests from an initiator to achieve a specified bandwidth. Choudhury does not control initiators at all, as pointed above, and it is respectfully submitted that the combination of Guha and Choudhury does not teach or suggest the recitations of Claims 6-8. Accordingly, it is respectfully submitted that the rejection of these Claims is improper and should be reversed.

3. No Combination of Guha and Choudhury Would Produce the Claimed Invention

Despite the fact that there is no teaching or suggestion in Guha and Choudhury to combine the references to meet the claimed invention, even if the references were combined as suggested by the Office, the combination would not produce the claimed invention. As pointed out previously, Guha does not disclose providing Quality of Service by a storage switch, but rather uses a separate device or appliance for this purpose. Accordingly, even assuming that Guha and Choudhury were combined, the combination would not produce a storage switch which provides Quality of Service, nor provide Quality of Service by controlling the number of packets or adjusting the number of requests from an initiator to a storage device, as claimed. Rather, the logical combination of Guha and Choudhury would be to incorporate the probability measurement techniques disclosed by Choudhury into the Quality of Service enforcer 34 of Guha to provide QoS to external requests to servers. Such a combination would not produce the invention in independent Claim 1 or in any of dependent Claims 3-8. Accordingly, the rejection of Claims 3-8 as obvious over the combination of Guha and Choudhury should be reversed.

E. Claims 9-15 Would Not Have Been Obvious Over Guha and Choudhury

Independent Claim 9 is directed to a method for use in a storage network that includes at least one initiator, at least one storage device, and at least one storage switch, where the initiator and storage device are in communication with the storage switch, and which the storage switch guarantees a minimum bandwidth to the initiator to access the storage device and measures the actual bandwidth utilized by the initiator, where the actual bandwidth is measured by a number of requests per second times an average size of a request.

The Office rejected independent Claim 9 on the basis that the features in the claim could also be found in Claims 1-8, and rejected Claim 9 for the same reasons as Claims 1 and 3-8 (see Office Action page 5, paragraph 12).

Claim 9 recites:

> A method for use in a storage network, the storage network including at least one initiator, at least one storage device, and at least one storage switch, wherein the at least one initiator and the at least one storage device are both in communication with the storage switch, the method comprising:
>
> guaranteeing, by the storage switch, a minimum bandwidth to the at least one initiator to access the at least one storage device in the storage network; and
>
> measuring, by the storage switch, an actual bandwidth utilized by the at least one initiator, where the actual bandwidth is measured by a number of requests per second times an average size of requests from the at least one initiator.

For the same reasons previously discussed in connection with the rejection of Claims 1 and 3-8, neither Guha nor Choudhury, alone or in combination, teaches or suggests a storage switch which performs a method of guaranteeing a minimum bandwidth to an initiator to access a storage device in a storage network, and measuring the actual bandwidth by measuring the number of requests per second times an average size of requests, as set forth in Claim 9. Moreover, there is no disclosure or suggestion in Guha or Choudhury individually or in combination of any element, much less a storage switch, performing a method comprising the steps set forth in Claim 9. Accordingly, the rejection of Claim 9 is improper and should be reversed.

Similarly, Claims 10-14 depend from Claim 9 and cannot be rendered obvious by the combination of the references for at least the same reasons Claim 9 cannot be rendered obvious. Accordingly, the rejection of Claims 10-14 is also improper and should be reversed.

F. Claims 16-21 and 38-44 Would Not Have Been Obvious Over The References For the Same Reasons Claims 9-15 Would Not Have Been Obvious

Independent Claim 16 is somewhat similar to Claim 9, and recites a method for use in a storage network that comprises:

> guaranteeing, by the storage switch, a respective minimum bandwidth for each of a plurality of connections, wherein each respective connection is a connection from a respective initiator to a respective target via the storage switch in the storage network;
>
> monitoring, by the storage switch, an actual bandwidth utilized by each initiator, where the actual bandwidth is measured by a number of requests

per second from the initiator times an average size of the requests from the initiator; and

determining if the actual bandwidth used by one initiator is excessive, and, if excessive, adjusting, by the storage switch, a number of allowed concurrent requests for at least one initiator.

Independent Claim 38 is directed to the machine readable media having instructions for execution by a storage switch to perform a method that comprises substantially the same steps of guaranteeing and measuring that are set forth in Claim 9.

Claim 16 recites a method that comprises guaranteeing and measuring steps substantially as set forth in Claim 9, and is deemed to be allowable for at least the same reasons discussed above with respect to Claim 9. Moreover, Claim 16 sets forth additionally determining if an actual bandwidth used by one initiator is excessive, and adjusting by the storage switch the number of allowed requests for the one initiator.

There is no disclosure in Guha or Choudhury of determining if an actual bandwidth used by an initiator is excessive, nor, as previously pointed out, does either reference disclose or suggest adjusting by a storage switch the number of allowed concurrent requests. Accordingly, Claim 16 is deemed allowable over the references for this reason also.

Claims 17-21, which depend from Claim 16, are deemed to be allowable for at least the same reasons.

Independent Claim 38 recites method steps that are substantially the same as recited in Claim 9, and Claim 38 is deemed to be allowable over the cited references for the same reasons as Claim 9 is allowable. Claims 39-44, which depend from Claim 38, are deemed to be allowable for the same reasons Claim 38 is allowable.

G. Claims 22-24 Would Not Have Been Obvious Over the Cited Prior Art

Independent Claim 22 sets forth a method for use in a storage network that includes an initiator, a storage device and a storage switch where the initiator and the storage device are in communication with the storage switch, and the method comprises providing a connection from the initiator to the storage device via the storage switch in the storage network, and adjusting by the storage switch the number of requests allowed the initiator to keep the bandwidth utilized by the initiator within a specified range.

For the same reasons discussed above in connection with Claims 1 and 6, the cited references to Guha and Choudhury do not disclose or suggest a storage switch which performs a method as set forth in Claim 22. Accordingly, it is respectfully submitted that Claim 22 is allowable over the cited prior art.

Claims 23-24, which depend from Claim 22, are deemed to be allowable for at least the same reasons Claim 22 is allowable.

H. Claims 27-29 and 31-32 Would Not Have Been Obvious Over the Cited Prior Art

Claims 27-29 depend from Claim 25, and Claims 31-32 depend from Claim 30, and these claims are deemed to be allowable over the cited prior art for at least the same reasons that their corresponding independent Claims 25 and 30 are allowable.

I. Claims 34-36 Are Patentable Over The Cited Preferences

Claims 34-36 depend from independent Claim 33. Claim 34, like Claim 33, is written in means-plus-function language. For the reasons discussed above in connection with the rejection of Claim 33, the Office has failed to properly construe Claim 34, since it has failed to construe the claim pursuant to 35 U.S.C. §112, ¶6 as corresponding to the structure disclosed in the specification for performing the recited functions, and equivalence thereof.

The specification discloses the structure for performing the recited functions as being a storage switch. Since Guha does not disclose a storage switch which performs the recited functions, for the reasons pointed out above, neither Guha alone or in combination with Choudhury teach or suggest the invention set forth in Claims 33-34. Accordingly, Claims 33-34 are deemed allowable over the cited prior art. Claims 35-36 depend from Claim 33 and are deemed allowable for at least the same reasons Claim 33 I s allowable.

VIII. CONCLUSION

In view of the foregoing, it is respectfully submitted that the rejections of Claims are improper, unsustainable, and should be reversed.

Dated: July 5, 2006

Respectfully Submitted,

Barry N. Young
Attorney for Assignee/Appellants
Reg. No. 27,744

Customer No. 48789
Law Offices of Barry N. Young
Court House Plaza, Suite 410
260 Sheridan Avenue
Palo Alto, CA 94306-2047
Phone: (650) 326-2701
Fax: (650) 326-2799
byoung@young-iplaw.com

A. CLAIMS APPENDIX

1. A method for use in a storage network, the storage network including at least one initiator, at least one storage device, and a storage switch in communication with the at least one initiator and the at least one storage device, the method comprising:

providing, by the storage switch, quality of service to the at least one initiator for accessing the at least one storage device in the storage network.

2. (cancelled)

3. The method of claim 1, wherein the step of providing quality of service includes controlling the number of packets from the at least one initiator to the at least one storage device during a period of time.

4. The method of claim 1, wherein the step of providing quality of service includes controlling the number of requests from the at least one initiator to the at least one storage device.

5. The method of claim 1, wherein the step of providing quality of service includes adjusting a number of concurrent requests allowed to be sent by the at least one initiator.

6. The method of claim 1, wherein the step of providing quality of service includes adjusting the number of requests allowed the at least one initiator to keep the bandwidth utilized by the at least one initiator within a specified range.

7. The method of claim 1, wherein the step of providing quality of service includes:

guaranteeing a minimum bandwidth to the at least one initiator to access the storage device;

measuring an actual bandwidth utilized by the initiator, where the actual bandwidth is measured by a number of requests per second times an average size of requests from the at least one initiator; and

adjusting a number of concurrent requests allowed to be sent by the at least one initiator.

8. The method of claim 7, further including:

guaranteeing up to a maximum bandwidth to the at least one initiator to access the at least one storage device;

wherein adjusting the number of concurrent requests includes reducing the number of concurrent requests allowed by the at least one initiator when the actual bandwidth exceeds the maximum bandwidth.

9. A method for use in a storage network, the storage network including at least one initiator, at least one storage device, and at least one storage switch, wherein the at least one initiator and the at least one storage device are both in communication with the storage switch, the method comprising:

guaranteeing, by the storage switch, a minimum bandwidth to the at least one initiator to access the at least one storage device in the storage network; and

measuring, by the storage switch, an actual bandwidth utilized by the at least one initiator, where the actual bandwidth is measured by a number of requests per second times an average size of requests from the at least one initiator.

10. The method of claim 9, further comprising:

adjusting a number of concurrent requests allowed to be sent by the at least one initiator.

11. The method of claim 10, wherein the step of adjusting includes:

reducing the number of concurrent requests allowed to be sent by the at least one initiator.

12. The method of claim 10, wherein the step of adjusting includes:

increasing the number of concurrent requests allowed to be sent by the at least one initiator.

13. The method of claim 9, further including guaranteeing, by the storage switch up to a maximum bandwidth to the at least one initiator to access the storage device.

14. The method of claim 13, further including:

reducing the number of concurrent requests allowed by the at least one initiator when the actual bandwidth exceeds its maximum bandwidth.

15. The method of claim 9, wherein measuring the actual bandwidth includes determining if a buffer includes a number of packets beyond a specified threshold.

16. A method for use in a storage network, the storage network including a plurality of initiators, a plurality of targets, and a storage switch, the method comprising:

guaranteeing, by the storage switch, a respective minimum bandwidth for each of a plurality of connections, wherein each respective connection is a connection from a respective initiator to a respective target via the storage switch in the storage network;

monitoring, by the storage switch, an actual bandwidth utilized by each initiator, where the actual bandwidth is measured by a number of requests per second from the initiator times an average size of the requests from the initiator; and

determining if the actual bandwidth used by one initiator is excessive, and, if excessive, adjusting, by the storage switch, a number of allowed concurrent requests for at least one initiator.

17. The method of claim 16, wherein monitoring the actual bandwidth includes determining if a buffer includes a number of packets beyond a specified threshold.

18. The method of claim 16, wherein adjusting a number of allowed concurrent requests includes reducing the number of allowed concurrent requests to the one initiator that is using excessive bandwidth.

19. The method of claim 18, wherein adjusting a number of allowed concurrent requests includes increasing the number of allowed concurrent requests to another initiator.

20. The method of claim 16, wherein the targets are virtual targets.

21. The method of claim 16, further including guaranteeing, by the storage switch, up to a respective maximum bandwidth for each of the plurality of connections, wherein determining if the actual bandwidth used by one initiator is excessive includes determining if the one initiator has exceeded its maximum bandwidth.

22. A method for use in a storage network, the storage network including at least one initiator, at least one storage device, and a storage switch, wherein the at least one initiator and the at least one storage device are both in communication with the storage switch, the method comprising:

providing a connection from the at least one initiator to the at least one storage device via the storage switch in the storage network; and

adjusting, by the storage switch, the number of requests allowed the at least one initiator to keep the bandwidth utilized by the at least one initiator within a specified range.

23. The method of claim 22, wherein bandwidth is defined by a number of requests per second from the at least one initiator times an average size of the requests from the at least one initiator.

24. The method of claim 22, wherein the number of requests allowed the at least one initiator is the number of concurrent requests allowed the at least one initiator.

25. A switch for use in a storage network, the switch comprising:

a port to be coupled to an external device, wherein the external device includes at least one of an initiator and a storage device; and

a bandwidth controller, the bandwidth controller including a processor, a traffic manager, and a buffer.

26. (cancelled)

27. The switch of claim 25, wherein the processor is a storage processor.

28. The switch of claim 25, wherein the port and the bandwidth controller are on one of a plurality of linecards in the switch, wherein each linecard includes a respective port and a respective bandwidth controller.

29. The switch of claim 25, wherein bandwidth is defined by a number of requests per second times an average size of the requests.

30. A switch, including:

a storage processor, including a request controller;

a traffic manager in communication with the storage processor;

a buffer in communication with the traffic manager;

wherein if a specified threshold in the buffer is reached, the traffic manager is designed to activate the request controller.

31. The switch of claim 30, wherein the request controller is designed to adjust the number of requests allowed an initiator to keep the bandwidth utilized by the initiator within a specified range.

32. The switch of claim 31, wherein bandwidth is defined by a number of requests per second times an average size of the requests.

33. A storage switch for use in a storage network comprising:

a first port to be coupled to at least one initiator;

a second port to be coupled to at least one storage device; and

means for providing quality of service for a connection from the at least one initiator to the at least one storage device in the storage network.

34. The switch of claim 33, wherein means for providing quality of service includes:

means for guaranteeing a minimum bandwidth to at least one initiator to access a storage device;

means for measuring an actual bandwidth utilized by the at least one initiator, where the actual bandwidth is measured by the number of requests per second times the average size of the requests from the at least one initiator; and

means for adjusting the number of concurrent requests allowed to be sent by the at least one initiator to keep the bandwidth utilized by the at least one initiator within a specified range having as a lower limit the minimum bandwidth.

35. The switch of claim 33, wherein means for providing quality of service includes:

a processor;

a traffic manager; and

a buffer.

36. The switch of claim 35, wherein the processor is a storage processor.

37. A storage network, including:

an initiator;

a storage device;

a switch in communication with the initiator and the storage device;

wherein the switch includes a traffic manager in communication with a buffer;

wherein when the buffer includes a number of packets from the initiator that exceeds a specified threshold, then the switch is designed to notify the initiator to reduce a number of concurrent requests.

38. A machine readable media which has instructions stored thereon, which when executed by a storage switch in a storage network including an initiator and a storage device in communication with the storage switch causes the storage switch to perform the following steps:

guaranteeing, by the storage switch, a minimum bandwidth to the *initiator* to access the storage device in the storage network; and

measuring, by the storage switch, an actual bandwidth utilized by the initiator, where the actual bandwidth is measured by a number of requests per second times an average size of requests from the initiator.

39. The machine readable media of claim 38, further including instructions for performing the step of:

adjusting a number of concurrent requests allowed to be sent by the initiator.

40. The machine readable media of claim 39, wherein the step of adjusting includes:

reducing the number of concurrent requests allowed to be sent by the initiator.

41. The machine readable media of claim 39, wherein the step of adjusting includes:

increasing the number of concurrent requests allowed to be sent by the initiator.

42. The machine readable media of claim 38, further including instructions for performing the step of:

guaranteeing, by the storage switch, up to a maximum bandwidth to the initiator to access the storage device.

43. The machine readable media of claim 42, further including instructions for performing the step of:

reducing the number of concurrent requests allowed by the initiator when it exceeds its maximum bandwidth.

44. The machine readable media of claim 38, wherein measuring the actual bandwidth includes determining if a buffer includes a number of packets beyond a specified threshold.

B. EVIDENCE APPENDIX

Exhibit A: Bird, D., "Network Storage- the Basics", Enterprise Storage Forum Technology Article, www.enterprisestorageforum.com/technology/features/articles.htm, January 2, 2002

Exhibit B: Bird, D., "Network Storage- Storage Area Networks", Enterprise Storage Forum Technology Article, www.enterprisestorageforum.com/technology/features/articles.htm, February 26, 2002

Exhibit C: Answers.com, definition of "Storage Switch", Information from Answers.com, www.answers.com/topic/storage-switch?hi=fibre&hi=ch . . .

Exhibit D: Answers.com, definition of "Fibre Channel Switch", Information from Answers.com, www.answers.com/topic/fibre-channel-switch?hi=storag . . .

www.enterprisestorageforum.com/technology/features/article.php/947551

Back to Article

Network Storage - The Basics
By Drew Bird
January 2, 2002

Are you new to network storage? If so then this series of articles is for you! Over the next few months we are going to take a look at the basic principles of network storage and answer questions like 'What is network storage?' and 'Why do we use it?' After covering the basics, subsequent articles will look at specific technologies in more detail. All of the articles in the series will have one simple aim; to educate and inform you about network storage. So, without further ado, lets get to it!

In basic terms, network storage is simply about storing data using a method by which it can be made available to clients on the network. Over the years, the storage of data has evolved through various phases. This evolution has been driven partly by the changing ways in which we use technology, and in part by the exponential increase in the volume of data we need to store. It has also been driven by new technologies, which allow us to store and manage data in a more effective manner.

In the days of mainframes, data was stored physically separate from the actual processing unit, but was still only accessible through the processing units. As PC based servers became more commonplace, storage devices went 'inside the box' or in external boxes that were connected directly to the system. Each of these approaches was valid in its time, but as our need to store increasing volumes of data and our need to make it more accessible grew, other alternatives were needed. Enter network storage.

Network storage is a generic term used to describe network based data storage, but there are many technologies within it which all go to make the magic happen. Here is a rundown of some of the basic terminology that you might happen across when reading about network storage.

Direct Attached Storage (DAS)
Direct attached storage is the term used to describe a storage device that is directly attached to a host system. The simplest example of DAS is the internal hard drive of a server computer, though storage devices housed in an external box come under this banner as well. DAS is still, by far, the most common method of storing data for computer systems. Over the years, though, new technologies have emerged which work, if you'll excuse the pun, out of the box.

Network Attached Storage (NAS)
Network Attached Storage, or NAS, is a data storage mechanism that uses special devices connected directly to the network media. These devices are assigned an IP address and can then be accessed by clients via a server that acts as a gateway to

EXHIBIT A

the data, or in some cases allows the device to be accessed directly by the clients without an intermediary.

The beauty of the NAS structure is that it means that in an environment with many servers running different operating systems, storage of data can be centralized, as can the security, management, and backup of the data. An increasing number of companies already make use of NAS technology, if only with devices such as CD-ROM towers (stand-alone boxes that contain multiple CD-ROM drives) that are connected directly to the network.

Some of the big advantages of NAS include the expandability; need more storage space, add another NAS device and expand the available storage. NAS also bring an extra level of fault tolerance to the network. In a DAS environment, a server going down means that the data that that server holds is no longer available. With NAS, the data is still available on the network and accessible by clients. Fault tolerant measures such as RAID, which we'll discuss later), can be used to make sure that the NAS device does not become a point of failure.

Storage Area Network (SAN)
A SAN is a network of storage devices that are connected to each other and to a server, or cluster of servers, which act as an access point to the SAN. In some configurations a SAN is also connected to the network. SAN's use special switches as a mechanism to connect the devices. These switches, which look a lot like a normal Ethernet networking switch, act as the connectivity point for SAN's. Making it possible for devices to communicate with each other on a separate network brings with it many advantages. Consider, for instance, the ability to back up every piece of data on your network without having to 'pollute' the standard network infrastructure with gigabytes of data. This is just one of the advantages of a SAN which is making it a popular choice with companies today, and is a reason why it is forecast to become the data storage technology of choice in the coming years. According to research company IDC, SAN's will account for 70% of all network storage by 2004.

Irrespective of whether the network storage mechanism is DAS, NAS or SAN, there are certain technologies that you'll find in almost every case. The technologies that we are referring to are things like SCSI and RAID. For years SCSI has been providing a high speed, reliable method for data storage. Over the years, SCSI has evolved through many standards to the point where it is now the storage technology of choice. Related, but not reliant on SCSI, is RAID. RAID (Redundant Array of Independent Disks) is a series of standards which provide improved performance and/or fault tolerance for disk failures. Such protection is necessary as disks account for 50% of all hardware device failures on server systems. Like SCSI, RAID, or the technologies used to implement it, have evolved, developed and matured over the years.

In addition to these mainstays of storage technology, other technologies feature in our network storage picture. One of the most significant of these technologies is Fibre channel (yes, that that's fiber with an 're'). Fibre Channel is a technology used to interconnect storage devices allowing them to communicate at very high speeds (up to 10Gbps in future implementations). As well as being faster than more traditional storage technologies like SCSI, Fibre Channel also allows for devices to be connected over a much greater distance. In fact, Fibre Channel can be used up to six miles. This allows devices in a SAN to be placed in the most appropriate physical location.

Other developments are coming through that will change the way that we use and access network storage. One of these advances pegged to make a large contribution to the growing success of network storage in general is iSCSI. iSCSI is a technology that allows data to be transported to and from storage devices over an IP network. What it actually does is serialize the data from a SCSI connection. Using iSCSI, the concept of network storage can be taken anywhere that IP can go, which as the Internet proves, is basically anywhere. Technologies like Fibre Channel and iSCSI are a big factor in how fast people are able to afford and implement network storage solutions.

Over the coming months, we'll be taking a detailed look at all of the technologies that we have discussed in this introductory article. In our next article we'll start by

taking a detailed look at perhaps the most significant element of today's network storage environment - SAN's. We'll also examine the devices used to create them. In addition, we'll be asking and answering the question 'How can a SAN benefit your business?' Stay tuned.

Click here to go to part II of Series: Storage Basics - Storage Area Networks

www.enterprisestorageforum.com/technology/features/article.php/981191

Back to Article

Storage Basics: Storage Area Networks
By Drew Bird
February 26, 2002

Many IT organizations today are scratching their heads debating whether the advantages of implementing a SAN solution justify the associated costs. Others are trying to get a handle on today's storage options and whether SAN is simply Network Attached Storage spelled backwards. In this article, I introduce the basic purpose and function of a SAN and examine its role in modern network environments. I also look at how SANs meet the network storage needs of today's organizations and answer the question, could a SAN be right for you.

Peel away the layers of even the most complex technologies and you are likely to find that they provide the most basic of functions. This is certainly true of storage area networks (SANs). Behind the acronyms and revolutionary headlines, lies a technology designed to provide a way of offering one of the oldest of network services, that of making access to data storage devices available to clients.

In very basic terms, a SAN can be anything from two servers on a network accessing a central pool of storage devices to several thousand servers accessing many millions of megabytes of storage. Conceptually, a SAN can be thought of as a separate network of storage devices physically removed from, but still connected to, the network. SANs evolved from the concept of taking storage devices, and therefore storage traffic, off the LAN and creating a separate back-end network designed specifically for data.

SANs represent the evolution of data storage technology to this point. Traditionally, on client server systems, data was stored on devices either inside or directly attached to the server. Next in the evolutionary scale came Network Attached Storage (NAS) which took the storage devices away from the server and connected them directly to the network. SANs take the principle one step further by allowing storage devices to exist on their own separate network and communicate directly with each other over very fast media. Users can gain access to these storage devices through server systems which are connected to both the LAN and the SAN.

This is in contrast to the use of a traditional LAN for providing a connection for server-storage, a strategy that limits overall network bandwidth. SANs address the bandwidth bottlenecks associated with LAN based server storage and the scalability limitations found with SCSI bus based implementations. SANs provide modular scalability, high-availability, increased fault tolerance and centralized storage management. These advantages have led to an increase in the popularity of SANs as they are quite simply better suited to address the data storage needs of today's data intensive network environments.

EXHIBIT B

The advantages of SANs are numerous, but perhaps one of the best examples is that of the serverless backup (also commonly referred to as 3rd Party Copying). This system allows a disk storage device to copy data directly to a backup device across the high-speed links of the SAN without any intervention from a server. Data is kept on the SAN, which means the transfer does not pollute the LAN, and the server processing resources are still available to client systems.

SANs are most commonly implemented using a technology called Fibre channel (yes, that's fibre with an 're', not an 'er'). Fibre Channel is a set of communication standards developed by the American National Standards Institute (ANSI). These standards define a high-performance data communications technology that supports very fast data rates (over 2Gbps). Fibre channel can be used in a point-to-point configuration between two devices, in a 'ring' type model known as an arbitrated loop, and in a fabric model.

Devices on the SAN are normally connected together through a special kind of switch, called a Fibre Channel switch, which performs basically the same function as a switch on an Ethernet network, in that it acts as a connectivity point for the devices. Because Fibre channel is a switched technology, it is able to provide a dedicated path between the devices in the fabric so that they can utilize the entire bandwidth for the duration of the communication.

The storage devices are connected to Fibre Channel switch using either multimode or single mode fiber optic cable. Multimode for short distances (up to 2 kilometers), single mode for longer. In the storage devices themselves, special fiber channel interfaces provide the connectivity points. These interfaces can take the form of built in adapters, which are commonly found in storage subsystems designed for SANs, or can be interface cards much like a network card, which are installed into server systems.

So, the question that remains is this. Should you be moving away from your current storage strategy and towards a SAN? The answer is not a simple one. If you have the need to centralize or streamline your data storage then a SAN may be right for you. There is, of course, one barrier between you and storage heaven, and that's money. While SANs remain the domain of big business, the price tag's of SAN equipment is likely to remain at a level outside the reach of small or even medium sized businesses. As the prices fall, however, SANs will find their way into organizations of all sizes, including, if you want, yours.

Click here to go to part I of Series: Network Storage - The Basics

storage switch

On this page:
Technology

Technology

Computer Desktop Encyclopedia

storage switch

A device that routes data between servers and disk arrays in a storage area network (SAN). It typically refers to a Fibre Channel switch. See SAN and Fibre Channel.

Mentioned In

storage switch is mentioned in the following topics:

Emulex	digital signal (Technology)
Fibre Channel switch	file protection (technology)
Snap Server (technology)	IP storage (technology)
KYK-13	Fibre Channel (technology)
Fibre channel port	LAN free backup (technology)
	More>

EXHIBIT C

Copyrights:

Computer Desktop Encyclopedia " Cite "

On this page: Technology E-mail Page Print this page Link to this page

Tell me about: storage switch Answers Web Shop

Home Webmasters Site Map About Help Advertise RSS

Preferences
Answers.com™ Fibre Channel switch Answers Web Shop
Business Entertainment Games Health People Places Reference Science Shopping Words
More...

Fibre Channel switch

On this page:
Wikipedia

Wikipedia

Fibre Channel switch

A Fibre Channel switch is a computer storage device that allows the creation of a Fibre Channel fabric. This fabric is a network of Fibre Channel devices which allows many-to-many communication, device name lookup, security, and redundancy. Major manufacturers of Fibre Channel switches are: Brocade, Cisco, IBM, McData and Qlogic

Fibre Channel fabrics are normally divided into zones to control access.

See also

- Host Truck Adapter (HTA)
- List of Fibre Channel switches
- List of Fibre Channel Host Bus Adapters
- Fibre Channel Port

Mentioned In

Fibre Channel switch is mentioned in the following topics:

storage switch (technology)	Director-class switch (technology)
Fibre Channel fabric	fabric (disambiguation)
List of Fibre Channel Host Bus Adapters	NAS gateway (technology)
LAN free backup (technology)	Brocade Communications Systems
SAN (technology)	IP storage (technology)
	More>

EXHIBIT D

Copyrights:

On this page: Wikipedia E-mail Page Print this page Link to this page

Tell me about: Fibre Channel switch Answers Web Shop

Home Webmasters Site Map About Help Advertise RSS